A Discussion of "All or None" Inspection Policies

Scott A. Vander Wiel
Department of Statistics and
Center of Nondestructive Evaluation
Iowa State University,
Ames, IA 50011

Stephen B. Vardeman
Department of Statistics,
Department of Industrial and Manufacturing Systems Engineering and
Center of Nondestructive Evaluation
Iowa State University,
Ames, IA 50011

August 31, 1992

Abstract

An optimal inspection policy will inspect either every item produced or no item when (1) product characteristics are well modeled as independent and identically distributed and (2) overall inspection cost is a sum of individually and identically determined costs for the items encountered. We show that even when independent inspection errors are possible, these conditions lead to optimality of either 0% or 100% inspection under several standard inspection scenarios and, more generally, even among fully sequential inspection plans. We also give examples to demonstrate that both "other" cost structures and "informative" inspections (that is, lack of independence) can lead to optimal policies which are *not* of the all or none type.

KEY WORDS: Acceptance sampling; Imperfect inspection; Cost optimality; Sequential decision making

1 Introduction

It is often said in product sampling inspection that all or none is optimal (e.g., Deming, 1986, Chapter 15). That is, in order to minimize cost, one should inspect (and, if defective, replace or rectify) either every item or no item from the production stream — the optimum choice depending on various costs involved and the expected fraction defective in the particular situation. In the production of critical metal components for aircraft engines it is generally agreed that 100% inspection is optimal because of the high cost (probably loss of lives) incurred when a defective component fails in service. However, the usual justification for all or none inspection does not include the possibility that inspection results may be imperfect. Typical nondestructive evaluation techniques using, for example, eddy currents or ultrasonics are associated with substantial measurement noise and this application has partly motivated our work.

In this paper we describe conditions under which all or none optimality holds and indicate some cases in which optimal policies need not be of the all or none type. We also discuss the issue of imperfect inspection and its effect on typical all or none criteria.

Optimality of all or none inspection has historically been discussed in the context of "go no–go" acceptance sampling of incoming lots of parts for subsequent assembly. See for example Mood (1943) and Barnard (1954) or, more recently, Deming (1982, Chapter 13) and Vardeman (1986). In Sections 2 through 4 we demonstrate that with appropriate assumptions optimal inspection policies are of the all or none type regardless of whether one is using fixed interval inspections, random inspections, the continuous sampling plan of Dodge (1943), or traditional lot acceptance sampling. In this way we broaden what has historically been a focus on all or none optimality for traditional acceptance sampling.

In Section 5 we show that under the following two conditions either 0% or 100% inspection is the optimal policy among a wide class of fully sequential inspection policies including those of Sections 2 through 4.

1. The physical characteristics of an item and the item's corresponding inspection results are well modeled with a joint probability distribution that is identical for each item and independent from item to item.

2. Overall costs are a sum of individual costs associated with an item's condition and/or inspection results, assessed identically and separately for each item of production.

We refer to a product stream and inspection procedure as stable if Condition 1 is satisfied. The perspective of Sections 2 through 5 is frequentist. In most of the paper we treat product quality and inspection results as if they were binary quantities. However, in Section 5 we show that under Conditions 1 and 2 all or none inspection is optimal much more generally.

In order to determine whether it is 0% or 100% inspection that is optimal under 1 and 2 one must use the common joint distribution of an item's physical characteristics and corresponding inspection results. In particular, for the binary case considered in Sections 2 through 4, one must use the probability p that an item is defective. If p is not known then the optimal policy is not realizable. One way to handle this difficulty is to explicitly deal with one's ignorance about p by assigning it a prior probability distribution. In this case Condition 1 is usually no longer satisfied and the optimal realizable policy need no longer have the all or none character. This type of Bayesian analysis is discussed in Sections 6 and 7. In Section 8 we briefly consider how a cost structure which does not treat each item separately according to Condition 2 can lead to other than all or none optimality. Section 9 gives some concluding remarks.

2 Random and Fixed Interval Inspections of "Stable" Product Streams

2.1 Perfect Inspection

The situation in which all or none optimality is probably most easily demonstrated is the following: Units produced by a certain manufacturing line are either good (coded $X_t = 0$, t denoting the serial number of an item) or defective ($X_t = 1$). Successive X_t's are modeled as statistically independent (Condition 1) with constant defective probability p; i.e., $X_t \sim$ iid Bernoulli(p). An inspection cost $k_1 > 0$ is assessed for any item that is inspected and an amount $k_2 > k_1$ is assessed for any defective item that is not inspected. (In this case cost is as described in Condition 2.) Each item is inspected with probability π and the choice about whether to inspect a given item is independent of any previous choices and the outcomes of any previous inspections. The goal is to choose π so as to minimize the (overall total cost or equivalently the) average cost per unit

$$\pi k_1 + (1 - \pi)pk_2 = pk_2 + \pi(k_1 - pk_2).$$

It is clear that this is minimized by taking $\pi = 0$ (no inspection) if $k_1 - pk_2 < 0$ (i.e., $p < k_1/k_2$) and by $\pi = 1$ (inspect all) otherwise. Notice that in order to implement mathematically optimal inspection in this case, *one must know the constant defective probability, p.*

2.2 Imperfect Inspection

A refinement of the above is to admit that, realistically, inspection cannot be performed perfectly, and thus allow misclassification probabilities

w_0 = probability of misclassifying a good item as defective, and

w_1 = probability of misclassifying a defective item as good.

The good/defective states X_t of the items are again modeled as iid Bernoulli(p) random variables and inspection results are modeled as independent given the actual states (good or defective) of the inspected units. Associated non-negative costs are assessed as follows:

k_1 = cost to inspect an item,

k_2 = cost when a defective item passes inspection or is not inspected,

k_3 = cost to attempt rectification of a good item deemed defective, and

k_4 = cost to rectify a defective item.

The per unit cost is then

$$\begin{aligned} &\pi[k_1 + k_2 w_1 p + k_3 w_0(1-p) + k_4(1-w_1)p] + (1-\pi)pk_2 \\ &= pk_2 + \pi\{(k_1 + w_0 k_3) - p[(1-w_1)(k_2 - k_4) + w_0 k_3]\}. \end{aligned}$$

This is minimized, over choice of π, by $\pi = 0$ (no inspection) if

$$p < p_{\mathrm{c}} \equiv \frac{k_1 + w_0 k_3}{(1-w_1)(k_2 - k_4) + w_0 k_3}$$

and by $\pi = 1$ (inspect all) otherwise. We should point out that although the optimal π is discontinuous in p, it is clear from the expression for per unit cost that if p is near p_{c}, the per unit cost changes very little as a function of π and hence all, none or anything in between are about equivalent in terms of expected cost (though obviously they differ in inspection load.) Also, the result shows that including the possibility of inspection errors does not affect the all or none nature of the solution. It does, however, push one in the direction of doing no inspection (i.e., where p_c is mapped into $(0,1)$, it is an increasing function both w_0 and w_1.)

Essentially the same analysis can be performed if, in the above situation, the inspection scheme is a "fixed interval" inspection plan in which every mth item is inspected. In this case the role of π is taken by $1/m$.

3 CSP-1 Inspection Plans and "Stable" Product Streams

Another common setup in which we can show that all or none inspection is optimal is that of the so called type 1 continuous sampling plan (CSP-1) introduced by Dodge (1943). Whetherill (1977) gives an overview of this plan including some important references. He also shows for the case of perfect inspection that all or none is optimal.

CSP-1 is an inspection plan conducted in cycles consisting of two phases: 100% inspection and fixed interval inspection. The plan begins in the 100% inspection phase in which all units are inspected until some number i (called the clearing interval) of consecutive units are judged not defective. Then, inspection switches to the sampling phase during which only every mth item is inspected. The sampling phase continues until a defective item is found; this completes a cycle and causes a return to the 100% inspection phase.

Under the probability model and cost structure of Section 2.2, it is possible to derive expressions for the mean cycle cost and the mean cycle length. As shown in Appendix A, minimizing the ratio of these means with respect to i and m results in all or none optimality depending on whether the average fraction defective p is larger or smaller than the critical value p_c.

4 Acceptance Sampling of Lots and iid Bernoulli Models

In this Section we consider optimality of all or none inspection in the historical context of acceptance sampling of incoming lots of parts to be assembled in a subsequent operation where lot disposal decisions are made based on counts of good and defective inspection results. See, for example, Deming (1986, Chapter 15). In this case a sample of n parts from a lot of N is inspected and parts judged to be defective in the sample are rectified. If more than some number c in the sample are judged defective, the remaining $N - n$ parts are inspected and apparently defective parts are rectified. Otherwise, if fewer than c are judged defective, the entire lot is accepted with no further inspection.

In this situation, the good/defective states of the N parts in a lot have traditionally been modeled as iid Bernoulli(p) variates as in Section 2.2. This is typically justified on one of two bases:

1. As appropriate where the lot can be thought of as derived from some "stable" production process having p as before.

2. As a good approximation to a finite population simple random sampling model (for $p =$ the *real* lot fraction defective) in cases where n/N is small.

Under the model and cost structure (k_1, k_2, k_3, k_4) of Section 2.2, the optimal acceptance sampling policy is

- if $p < p_c$, then inspect $n = 0$ and always accept the lot, and
- if $p > p_c$, then inspect $n = N$.

See Appendix B for details. That is, all or none is optimal and the critical cost ratio is again the same as for the random inspection scenario of Section 2.2.

5 Unifying Comments on the Three Inspection Scenarios

The random inspection plan and the fixed interval plan might be termed non-sequential inspection plans since all of the units to be inspected can be determined before any of the inspection results are known. On the other hand, CSP-1 and the traditional acceptance sampling plan must be carried out sequentially. Decisions must be made "on the run" regarding which units are to be inspected.

The most general sequential sampling plan would allow for a decision on how to proceed further after each inspection result becomes available. Under the fixed p independence model for the good/defective states of units and our description of inspection efficacy in terms of w_0 and w_1, it is fairly straightforward to use dynamic programming to show that the optimal general sequential inspection policy depending on the apparently good/apparently defective inspection results where N parts are involved, is to do no inspection if $p < p_c$ and inspect everything otherwise. In Appendix C we show that all or none optimality holds even in a more general context where product characteristics and inspection measurements are possibly continuous vector valued random variables and Conditions 1 and 2 are satisfied.

Qualitatively, the reason that all or none holds so universally under Conditions 1 and 2 is that any information gained on any number of units has no effect on one's assessment of the likely characteristics of the remaining units. Hence, allowing sequential decision making is of no extra value. ¿From this point of view it should be no surprise that costs in all of the scenarios considered thus far are optimized by use of the same all or none rule.

6 More Complex Models

In this and the remaining sections we return to the case where product characteristics and inspection results are treated (possibly after classification based on continuous quantities) as binary quantities. Our statement in Section 1 of the two conditions under which all or none inspection is optimal suggests that it may be reasonable to inspect a fraction of the items if doing so can improve one's information about the condition of the remaining items. This is true, and can be understood by considering the extreme case of a manufacturer which ships perfect lots of parts yet occasionally makes shipments of the wrong stock number (i.e., 100% defective lots). If inspection is not too costly yet there is a cost associated with "defective" parts being passed on to assembly, then it is obvious that an optimal plan is to inspect one item from the lot to determine whether the lot is entirely good or entirely defective and replace the whole lot if needed, or route it to assembly otherwise.

More generally, one may wish to describe imprecise knowledge of p, or a time varying p (successive values of which we denote as $p_1, p_2, \ldots$), by means of a joint probability structure for $p_1, p_2, \ldots$. Doing so can lead to many varied and interesting problems in sequential decision making. We will here describe two very narrowly defined problems of this type for which solutions are straightforward.

First we note that it is possible to apparently generalize any of the fixed p situations described in Sections 2 through 4, by adding the possibility that each unit produced has a propensity to be defective, p_t, distributed independently according to a probability density $g(p)$ with mean p_0. (The states of the units are then modeled as independent given the defective probabilities $\{p_t\}$.) Such a model might be contemplated, for example, if before the manufacture of each part a machine were loaded with a cutting tool drawn randomly from a bin of tools. However, this probability structure for the states is really no different from that obtained when each item has a fixed defective probability p_0, which except for a slight notational change, is the case we have already examined.

On the other hand, if we model all units as having a common but unknown defective probability ($p_1 = p_2 = \cdots = p$ where p is described by a probability distribution $g(p)$), then it is possible to improve one's knowledge of uninspected units by knowing the fraction of apparently defective items found among those inspected. In these cases optimal inspection plans need not be of the all or none type and globally optimal plans will generally be of a fully sequential nature. In the next section we will discuss the problem of finding optimum "single sampling" plans in this context.

A much more complex situation is met if p_t is modeled as varying unit to unit according to a serially dependent stochastic process, like for example a Markov process. We know of no clean statement and solution to a nontrivial inspection problem of this more general type.

7 Bayesian Acceptance Sampling

If p, the probability of being defective in the acceptance sampling setup described in Section 4, is the same for each item and unknown but described with a probability or density function $g(p)$, the optimal initial inspection sample size n, need not be 0 or N (unless g puts all of its probability to one side of the critical cost ratio, in which case either all or none is optimal depending on where the probability is located). A "single sampling with rectification" acceptance sampling rule is a particular kind of two-stage decision rule. An initial sample is inspected, defective units are rectified, and the remainder of the lot is inspected and rectified only if the initial sample had more than some number c of defective units. Thyregod (1974) develops optimal acceptance sampling plans of this type in a general context allowing the state of each part to be a measurement on a continuous scale and individual item costs are charged correspondingly. Lorenzen (1985) specializes Thyregod's arguments to the Bernoulli situation.

For convenient densities $g(p)$, it is possible, with the aid of computer programs, to find values of n and c which minimize the expected total cost associated with inspection. Lorenzen's (1985) program e.g., uses a Beta density for $g(p)$ and assumes perfect inspection. Appendix D shows that the acceptance sampling problem incorporating possible inspection errors is mathematically the same as the one discussed by Lorenzen, aside from a rescaling of the prior distribution and a transformation of the cost parameters. However, we must point out that among all possible *fully sequential* inspection policies, even the best of these *two stage* rules is not generally globally optimal.

8 Other Overall Cost Structures

To understand the importance of Condition 2 to the all or none conclusions of this paper, consider a scenario in which a single working component out of an early shipment of several such components will later be needed to complete a high-priority project (e.g., a space shuttle). Suppose that inspection of the components can be done upon receipt at a moderate per item cost, and should the lot fail to contain at least one functioning component it can be returned to the supplier for replacement at minimal cost. On the other hand, if no incoming testing is done and it turns out at the time of attempted installation that the lot contains no good components, a huge monetary penalty associated with project delay will be suffered. Even if the Bernoulli model is appropriate and p is known and large, an optimal incoming inspection policy will *sample* until the first good component in the lot is found. Mathematization of this scenario involves a cost structure which does not treat each item separately as required by Condition 2 and produces something other than an all or none optimal rule.

9 Conclusion

Under a cost structure which assesses costs as a total of separately determined costs for each item of production and when information on the condition of some parts in a product stream cannot alter one's assessment of the likely condition of other parts, the best inspection policy is either to inspect every item or to let all items pass with no inspection. We have shown this to be true under several scenarios. Cases we considered are (1) random or fixed interval inspections, (2) inspection under a continuous sampling plan, and (3) acceptance sampling of lots. More generally, we have shown that the optimal fully sequential inspection plan for a stable product stream when costs are assessed separately and identically for each item is of the all or none variety, with the choice depending on the joint distribution of product characteristics and inspection results and the particular cost structure being used. In all of these arguments, we have allowed for the possibility that inspection results may not perfectly reflect the states of the inspected parts.

If the appropriate inspection cost structure is not a sum of separately and identically determined component costs, or if by inspecting some parts one can obtain an improved assessment of the remaining parts (which is the situation for nontrivial Bayesian models) then optimal inspection rules need not be of the all or none type. We have shown that adding the possibility of inspection errors to the Thyrefod/Lorenzen "single sampling with rectification" formulation of this problem leaves it unchanged except for a rescaling of the prior distribution on the probability of judging a part defective and a transformation of the cost parameters.

10 Acknowledgement

We wish to express our thanks to William Meeker for many useful suggestions on the content and presentation of this paper.

11 References

Barnard, G. A. (1954). "Sampling Inspection and Statistical Decisions." *Journal of the Royal Statistical Society* B16, pp. 151-174.

Deming, W. E. (1982). *Quality Productivity and Competitive Position.* MIT Center for Advanced Engineering Study, Cambridge, MA.

Deming, W. E. (1986). *Out of the Crisis.* MIT Center for Advanced Engineering Study, Cambridge, MA.

Dodge, H. F. (1943). "A Sampling Inspection Plan for Continuous Production." *Annals of Mathematical Statistics* 14, pp. 264-279.

Lorenzen, T. J. (1985). "Minimum Cost Sampling Plans Using Bayesian Methods." *Naval Research Logistics Quarterly* 32, pp. 57–69.

Mood, A. M. (1943). "On the Dependence of Sampling Inspection Plans Upon the Population Distributions." *The Annals of Mathematical Statistics* 14, pp. 415–425.

Thyregod, P. (1974). "Bayesian Single Sampling Acceptance Plans for Finite Lot Sizes." *Journal of the Royal Statistical Society* B36, pp. 305–319.

Vardeman, S. B. (1986). "The Legitimate Role of Inspection in Modern SQC." *The American Statistician* 40, pp. 325–328.

Wetherill, G. B. (1977). *Sampling Inspection and Quality Control.* Chapman and Hall, London.

Appendix A All or None Under CSP-1

Below, we minimize the ratio of mean cycle cost to mean cycle length under a CSP-1 inspection plan to show that the optimal plan has the all or none character.

With fixed p and under the imperfect inspection model set out in Section 2, the 100% inspection phase of a CSP-1 cycle has mean length given by

$$L_1 = \frac{1-(1-v)^i}{v(1-v)^i}$$

where $v \equiv p(1-w_1) + (1-p)w_0$ is the probability of (rightly or wrongly) judging an item to be defective. The average cost associated with a 100% inspection phase is

$$C_1 = \frac{1-(1-v)^i}{v(1-v)^i}[k_1 + k_2 w_1 p + k_3 w_0 (1-p) + k_4 (1-w_1)p].$$

The sampling phase of a CSP-1 cycle has a mean length of $L_2 = m/v$ and mean cost

$$C_2 = \frac{m}{v}\left[\frac{k_1}{m} + \frac{k_2 w_1 p}{m} + k_2\left(1-\frac{1}{m}\right)p + \frac{k_3 w_0 (1-p)}{m} + \frac{k_4(1-w_1)p}{m}\right].$$

The ratio of mean cycle cost to mean cycle length is (with a little algebra)

$$\begin{aligned} R &= \frac{C_1 + C_2}{L_1 + L_2} \\ &= k_1 + k_2 w_1 p + k_3 w_0 (1-p) + k_4 (1-w_1) p \\ &\quad + \frac{-\{k_1 + k_3 w_0 - p[(1-w_1)(k_2-k_4) + w_0 k_3]\}[(m-1)(1-v)^i]}{1+(m-1)(1-v)^i}. \end{aligned}$$

Viewed as a function of $x = (m-1)(1-v)^i$, this is of the form

$$R(x) = A + \frac{Bx}{1+x}$$

where A and B are constants. R is minimized over nonnegative x by $x = 0$ if $B > 0$ and by $x = \infty$ if $B < 0$. This corresponds to taking $m = 1$ (inspect all) if $p > p_c$ and $m = \infty$ (inspect none) if $p < p_c$.

Appendix B All or None Under Acceptance Sampling of Lots

Here we show that all or none is optimal when doing acceptance sampling of lots as described in Section 4. The expected cost associated with inspecting a lot is

$$\begin{aligned} C &= n[k_1 + k_2 w_1 p + k_3 w_0 (1-p) + k_4(1-w_1)p] \\ &\quad +(N-n)\{(1-P_{\mathrm{a}})[k_1 + k_2 w_1 p + k_3 w_0(1-p) + k_4(1-w_1)p] + P_{\mathrm{a}} k_2 p\} \\ &= N[k_1 + k_2 w_1 p + k_3 w_0(1-p) + k_4(1-w_1)p] \\ &\quad -P_{\mathrm{a}}(N-n)\{(k_1 + w_0 k_3) - p[(1-w_1)(k_2-k_4) + w_0 k_3]\} \end{aligned}$$

where P_{a} is the probability of accepting the lot under a given rule for lot disposal. If the factor in braces is positive, (i.e., $p < p_{\mathrm{c}}$) an optimal policy will have $P_{\mathrm{a}} = 1$ and $n = 0$; that is, accept the lot with no inspection. On the other hand if $p > p_{\mathrm{c}}$, then an optimal policy will have $P_{\mathrm{a}}(N-n) = 0$; that is, an optimal policy will examine every item in the lot.

Appendix C All or None and General Sequential Inspection Plans

Here we use backwards induction to show that under Conditions 1 and 2 an all or none policy is optimal among all possible sequential inspection policies, for a run of N items, that depend upon possibly noisy inspection results. Let X_t be a possibly vector valued random variable denoting the physical characteristics of the tth unit and let Y_t be a possibly vector valued random variable denoting the inspection results of the tth unit. Condition 1 states that the (X_t, Y_t) vectors are jointly iid. Let the inspection policy be given by the sequence $\{u_t\}$ with $u_t = 1$ indicating that the tth item is to be inspected and $u_t = 0$ otherwise. Each u_t is allowed to depend on inspection outcomes obtained before the tth item is presented for possible inspection. That is, u_t is potentially a function of the previous inspection choices $u_1, \ldots, u_{t-1}$ and the variables

$$Z_i = u_i Y_i, \qquad i = 1, 2, \ldots, t-1.$$

We seek an inspection policy to minimize the appropriate expected cost criterion which is

$$C = E\left\{\sum_{t=1}^{N}[u_t S(X_t, Y_t) + (1-u_t)A(X_t)]\right\}$$

where

$S(X,Y)$ denotes the cost of sampling, inspecting and either rejecting or accepting for use an item with characteristics X and inspection measurements Y; and

$A(X)$ denotes the cost of accepting for use with no inspection an item with characteristics X.

We note that one could also allow A to depend on Y_t but this seems unnecessary since the measurements Y_t are taken only if the item is sampled.

Below we will use the expectations

$$\begin{aligned} E[S(X_t,Y_t)|Z_1,\ldots,Z_{t-1}] &= E[S(X_t,Y_t)] \equiv \mathcal{S} \\ E[A(X_t)|Z_1,\ldots,Z_{t-1}] &= E[A(X_t)] \equiv \mathcal{A}. \end{aligned}$$

Let V_t^* be the optimal cost obtainable from item t to item N. Then

$$\begin{aligned} V_N^* &= \min_{u_N\in\{0,1\}} E\{u_N S(X_N,Y_N) + (1-u_N)A(X_N)\} \\ &= \min\{\mathcal{S},\mathcal{A}\} \end{aligned}$$

Hence, the optimal final period decision is $u_N^* = 1$ (inspect item N) if $\mathcal{S} < \mathcal{A}$ and $u_N^* = 0$ otherwise. Note that u_N^* and V_N^* do not depend on previous inspection decisions $u_1,\ldots,u_{N-1}$ or on the observed history $Z_1,\ldots,Z_{N-1}$.

Now we use induction. Suppose

$$V_{t+1}^* = (N-t)\min\{\mathcal{S},\mathcal{A}\}$$

Then

$$\begin{aligned} V_t^* &= \min_{u_t\in\{0,1\}} E\{u_t S(X_t,Y_t) + (1-u_t)A(X_t) + V_{t+1}^* \mid Z_1,\ldots,Z_{t-1}\} \\ &= \min_{u_t\in\{0,1\}}\{u_t\mathcal{S} + (1-u_t)\mathcal{A} + (N-t)\min\{\mathcal{S},\mathcal{A}\}\} \\ &= [N-(t-1)]\min\{\mathcal{S},\mathcal{A}\}. \end{aligned}$$

Hence, the optimal decision in period t is

$$\begin{aligned} u_t^* &= 1, \quad \text{if } \mathcal{S} < \mathcal{A} \\ &= 0, \quad \text{otherwise.} \end{aligned}$$

We see, in general, that u_t^* and V_t^* do not depend on previous inspection decisions $u_1, \ldots, u_{t-1}$ or on the observed history $Z_1, \ldots, Z_{t-1}$. Thus, the optimal (potentially sequential) inspection policy is to inspect all items if $\mathcal{S} < \mathcal{A}$ or no items otherwise.

The above argument holds for any fixed run size N and, since the solution is independent of N, it may be regarded as a solution to the infinite horizon problem of inspecting a never-ending stream of items.

For the Bernoulli case considered throughout most of the paper we have that (X_t, Y_t) are iid with $X_t \sim \text{Bernoulli}(p)$ and $Y_t \sim$ Bernoulli with

$$\begin{aligned} P[Y_t = 1] &= w_0 \quad && \text{if} \quad X_t = 0 \\ &= 1 - w_1 \quad && \text{if} \quad X_t = 1 \end{aligned}$$

and we would take the cost functions to be

$$S(X, Y) = k_1 + k_2 X(1 - Y) + k_3(1 - X)Y + k_4 XY$$

and

$$A(X) = k_2 X.$$

In this setup we have

$$\mathcal{S} = k_1 + k_2 p w_1 + k_3(1 - p) w_0 + k_4 p(1 - w_1)$$

and

$$\mathcal{A} = k_2 p$$

and the condition $\mathcal{S} < \mathcal{A}$ is equivalent to $p > p_c$.

Appendix D Bayesian Acceptance Sampling and Inspection

We here show that including the possibility of inspection errors in the Bayesian acceptance sampling problem of Lorenzen (1985), can be thought of in terms of Lorenzen's framework with a "rescaled" prior distribution and transformed cost parameters.

Lorenzen (1985) considered a perfect inspection model ($w_0 = w_1 = 0$) where, for given p, the elements of $\mathbf{X}_N \equiv (X_1, \ldots, X_N)$ are independent identically distributed Bernoulli variates with $P[X_t = 1] = p$ and the uncertainty in p is described using a Beta prior distribution. He considered three types of inspection costs, each depending on the condition of the part, namely: cost to sample a part, $S(X)$; cost to reject an unsampled part, $R(X)$; and cost to accept an unsampled part $A(X)$. For notational simplicity let

$$\begin{aligned} S(X) &= s_0, \quad \text{if } X = 0 \\ &= s_1, \quad \text{if } X = 1 \end{aligned}$$

$$\begin{aligned} R(X) &= r_0, \quad \text{if } X = 0 \\ &= r_1, \quad \text{if } X = 1 \end{aligned}$$

$$\begin{aligned} A(X) &= a_0, \quad \text{if } X = 0 \\ &= a_1, \quad \text{if } X = 1. \end{aligned}$$

Lorenzen's total cost function is

$$C(\mathbf{X}_N; n, A) = \sum_{t=1}^{n} S(X_t) + \sum_{t=n+1}^{N} R(X_t) + I[\mathbf{X}_n \in A] \sum_{t=n+1}^{N} (A(X_t) - R(X_t))$$

where $A \subset \{0,1\}^n$ is the acceptance region defining the decision rule. The problem is to minimize $E[C]$ with respect to both n and A. Thus, it suffices to consider minimization of the mean of a smoothed cost function with the same expectation as C:

$$\begin{aligned} C_1(\mathbf{X}_n, p; n, A) &= E\left[\sum_{t=1}^{n} S(X_t) \middle| p\right] + E\left[\sum_{t=n+1}^{N} R(X_t) \middle| p\right] \\ &\quad + E\left[I[\mathbf{X}_n \in A] \sum_{t=n+1}^{N} (A(X_t) - R(X_t)) \middle| \mathbf{X}_n, p\right] \\ &= n[s_0 + p(s_1 - s_0)] + (N - n)[r_0 + p(r_1 - r_0)] \\ &\quad + (N - n) I[\mathbf{X}_n \in A][a_0 + (p(a_1 - a_0) - r_0 - p(r_1 - r_0)]. \end{aligned}$$

We generalize Lorenzen's version of this problem as follows: $\mathbf{X}_N$ and p are modeled as above; however, the inspection results $\mathbf{Y}_N \equiv (Y_1, \ldots, Y_N)$

($Y_t = 1$ indicates the tth item is judged defective) are not necessarily the same as the states $\mathbf{X}_N$. Conditional on $\mathbf{X}_N$ and p, the elements of $\mathbf{Y}_N$ are modeled as independent Bernoulli variates with

$$\begin{aligned} P[Y_t = 1] &= w_0 \quad &\text{if} \quad X_t = 0 \\ &= 1 - w_1 \quad &\text{if} \quad X_t = 1. \end{aligned}$$

This implies that given $p_y \equiv (1-p)w_0 + p(1-w_1)$, the elements of $\mathbf{Y}_n$ are independent Bernoulli(p_y) variates. The distribution of p_y implied by the standard Beta prior on p is a Beta distribution scaled to the interval $(w_0, 1-w_1)$. (For convenience we assume $w_0 < 1 - w_1$.)

Our generalization of Lorenzen's cost function for the inspection error problem is

$$\begin{aligned} C'(\mathbf{X}_N, \mathbf{Y}_N; n, A) &= \sum_{t=1}^{n} S(X_t, Y_t) + \sum_{t=n+1}^{N} R(X_t, Y_t) \\ &\quad + I[\mathbf{Y}_n \in A] \sum_{t=n+1}^{N} (A(X_t) - R(X_t, Y_t)) \end{aligned}$$

where

$$\begin{aligned} S(X,Y) &= s_{00} \quad \text{if} \quad X = 0, Y = 0 \\ &= s_{01} \quad \text{if} \quad X = 0, Y = 1 \\ &= s_{10} \quad \text{if} \quad X = 1, Y = 0 \\ &= s_{11} \quad \text{if} \quad X = 1, Y = 1 \end{aligned}$$

$$\begin{aligned} R(X,Y) &= r_{00} \quad \text{if} \quad X = 0, Y = 0 \\ &= r_{01} \quad \text{if} \quad X = 0, Y = 1 \\ &= r_{10} \quad \text{if} \quad X = 1, Y = 0 \\ &= r_{11} \quad \text{if} \quad X = 1, Y = 1 \end{aligned}$$

$$\begin{aligned} A(X) &= \alpha_0, \quad \text{if} \quad X = 0 \\ &= \alpha_1, \quad \text{if} \quad X = 1. \end{aligned}$$

C' is actually more general than required for the cost structure of Sections 2 and 4 involving k_1, k_2, k_3 and k_4. To see that our costs can be written in Lorenzen's form, one uses the correspondences

$$\begin{aligned} s_{00} = r_{00} &= k_1, \\ s_{01} = r_{01} &= k_1 + k_3, \\ s_{10} = r_{10} &= k_1 + k_2, \end{aligned}$$

$$
\begin{aligned}
s_{11} = r_{11} &= k_1 + k_4, \\
\alpha_0 &= 0, \\
\alpha_1 &= k_2.
\end{aligned}
$$

Consider a smoothed version of C':

$$
\begin{aligned}
C_1'(\mathbf{Y}_n, p_y; n, A) &= E\left[\sum_{t=1}^{n} S(X_t, Y_t)\middle| p_y\right] + E\left[\sum_{t=n+1}^{N} R(X_t, Y_t)\middle| p_y\right] \\
&\quad + E\left[I[\mathbf{Y}_n \in A] \sum_{t=n+1}^{N} (A(X_t) - R(X_t, Y_t))\middle| p_y, \mathbf{Y}_n\right].
\end{aligned}
$$

We show below that C_1' is of the same form as C_1. Also, apart from the priors on p and p_y, $(\mathbf{X}_n, p)$ has the same probability structure as $(\mathbf{Y}_n, p_y)$. Hence, except for rescaling of the prior, and transforming the cost parameters, the problem with inspection errors is identical to the one with no errors.

To show that C_1' has the same form as C_1 we use $p_y = w_0 + p(1 - w_0 - w_1)$ to write

$$
\begin{aligned}
E[S(X_t, Y_t) \mid p_y] &= [(1-w_0)s_{00} + w_0 s_{01}] + p\{[(1-w_1)s_{11} + w_1 s_{10}] - [(1-w_0)s_{00} + w_0 s_{01}]\} \\
&= \{(1-w_1)[(1-w_0)s_{00} + w_0 s_{01}] - w_0[(1-w_1)s_{11} + w_1 s_{10}] \\
&\quad + p_y\{[(1-w_1)s_{11} + w_1 s_{10}] - [(1-w_0)s_{00} + w_0 s_{01}]\}\}/(1 - w_0 - w_1),
\end{aligned}
$$

$$
\begin{aligned}
E[R(X_t, Y_t) \mid p_y] &= [(1-w_0)r_{00} + w_0 r_{01}] + p\{[(1-w_1)r_{11} + w_1 r_{10}] - [(1-w_0)r_{00} + w_0 r_{01}]\} \\
&= \{(1-w_1)[(1-w_0)r_{00} + w_0 r_{01}] - w_0[(1-w_1)r_{11} + w_1 r_{10}] \\
&\quad + p_y\{[(1-w_1)r_{11} + w_1 r_{10}] - [(1-w_0)r_{00} + w_0 r_{01}]\}\}/(1 - w_0 - w_1).
\end{aligned}
$$

For $t > n$, we have

$$
\begin{aligned}
E[A(X_t) \mid \mathbf{Y}_n, p_y] &= (1-p)\alpha_0 + p\alpha_1 \\
&= [(1-w_1)\alpha_0 - w_0\alpha_1 + p_y(\alpha_1 - \alpha_0)]/(1 - w_0 - w_1)
\end{aligned}
$$

and

$$
E[R(X_t, Y_t) \mid \mathbf{Y}_n, p_y] = E[R(X_t, Y_t) \mid p_y].
$$

Hence,

$$
\begin{aligned}
C_1'(\mathbf{Y}_n, p_y; n, A) &= n\{(1-w_1)[(1-w_0)s_{00} + w_0 s_{01}] - w_0[(1-w_1)s_{11} + w_1 s_{10}] \\
&\quad + p_y\{[(1-w_1)s_{11} + w_1 s_{10}] - [(1-w_0)s_{00} + w_0 s_{01}]\}\}/(1 - w_0 - w_1)
\end{aligned}
$$

$$
\begin{aligned}
&+N\{(1-w_1)[(1-w_0)r_{00}+w_0r_{01}]-w_0[(1-w_1)r_{11}+w_1r_{10}]\\
&+p_y\{[(1-w_1)r_{11}+w_1r_{10}]-[(1-w_0)r_{00}+w_0r_{01}]\}\}/(1-w_0-w_1)\\
&+(N-n)I[\mathbf{Y}_n\in A]\times\{[(1-w_1)\alpha_0-w_0\alpha_1+p_y(\alpha_1-\alpha_0)]\\
&-(1-w_1)[(1-w_0)r_{00}+w_0r_{01}]-w_0[(1-w_1)r_{11}+w_1r_{10}]\\
&+p_y\{[(1-w_1)r_{11}+w_1r_{10}]-[(1-w_0)r_{00}+w_0r_{01}]\}\}/(1-w_0-w_1).
\end{aligned}
$$

This is in the form of C_1 with

$$
\begin{aligned}
s_0 &= \{(1-w_1)[(1-w_0)s_{00}+w_0s_{01}]-w_0[(1-w_1)s_{11}+w_1s_{10}]\}/(1-w_0-w_1)\\
s_1 &= \{(1-w_0)[(1-w_1)s_{11}+w_1s_{10}]-w_1[(1-w_0)s_{00}+w_0s_{01}]\}/(1-w_0-w_1)\\
r_0 &= \{(1-w_1)[(1-w_0)r_{00}+w_0r_{01}]-w_0[(1-w_1)r_{11}+w_1r_{10}]\}/(1-w_0-w_1)\\
r_1 &= \{(1-w_0)[(1-w_1)r_{11}+w_1r_{10}]-w_1[(1-w_0)r_{00}+w_0r_{01}]\}/(1-w_0-w_1)\\
a_0 &= [(1-w_1)\alpha_0-w_0\alpha_1]/(1-w_0-w_1)\\
a_1 &= [(1-w_0)\alpha_1-w_1\alpha_0]/(1-w_0-w_1)
\end{aligned}
$$

Since $\sum Y_t$ is sufficient for p_y and the distributions of Y_t and $\sum Y_t$ have monotone likelihood ratio, then the algorithm of Lorenzen (1985, Section 3) can be used to find optimal n and c. For the case with no inspection errors Lorenzen gives a computer program implementation which makes use of simplifications arising because the Beta family is conjugate for the Bernoulli model. For the case with inspection errors these simplifications are not available and implementing the algorithm would be more difficult.